

Fixed-Length Packets versus Variable-Length Packets in Fast Packet Switching Networks

Andrew Shaw

3 March 1994

Abstract

Fast Packet Switching (FPS) networks are designed to carry many kinds of traffic, including voice, video, and data. In this paper, we evaluate one design parameter of FPS networks: whether the packets should be fixed-length or variable-length. We consider three measures of performance: user frame loss rates, average latency, and effective bandwidth.

1 Introduction

This paper examines the technical merits of fixed-sized packets versus variable-sized packets for Fast Packet Switching (FPS) networks.

In order to understand the demands required of FPS networks, we give a short introduction to historical and technical issues which motivated the consensus about FPS as the appropriate architecture for an Integrated Services Network. Those who are familiar with Fast Packet Switching networks may choose to skip directly to the next section.

1.1 Convergence: Broadband Integrated Services Digital Network

Today, there are many different telecommunications networks which are operating to give many different kinds of service. Among these networks are the following [2] [9]:

- The telephone system – most of the traffic is point-to-point, and consists of voice, although teleconferencing, fax, and computer modem traffic is also increasing rapidly.
- TV – television is transported primarily by local-area broadcast, cable (Community Access TV), or satellite.
- Telex network – the Telex network transports character messages at a very slow rate (up to 300 bit/s).
- Local Area Networks (LANs) – computer systems within a single domain (i.e. company or school) are generally connected with relatively low bandwidth networks such as Ethernet, token bus, or token ring.
- Wide Area Networks (WANs) – computer systems which are geographically distributed are generally connected by a packet switched data networks. The most wide-reaching of these networks is the Internet.

It is interesting to note the demise of Telex network as a result of the general availability of fax machines. Although the original intent of these two networks was quite different, because of the greater availability of the telephone system and the greater flexibility of the fax (in transmitting images as well as text), the telephone system has absorbed the functionality of the Telex.

Although the demands of the various applications currently implemented on different networks may be quite different, there are many advantages to *integrating* the services into one network. One physical advantage is having one wire carry all the traffic to each customer (instead of having separate wires for telephone, data, and cable); another advantage is the prospect of universal service, which creates new opportunities for services which would be uneconomical or limited if implemented on a sparsely distributed network (which was one of the main reasons for the demise of the Telex).

These advantages led to the development of the Integrated Services Digital Network (ISDN) standard, which would provide the single wire, as well as the universal service. Unfortunately, as is described in [17] by Turner, the architecture for the ISDN network was essentially two completely different networks – one for voice, one for data – which would only be *physically* packaged in the same box and wire. In the same paper [17], Turner sketches out an architecture for a unified network which could handle both kinds of traffic, and which was the basis for the standards being considered today for Fast Packet Switching.

1.2 Packet Switching vs. Circuit Switching

Because the demands for voice and data traffic are so different, the architectures for voice networks (i.e. the phone system) and data traffic (i.e. LANs and WANs) were originally quite different. For voice traffic, the delay of transmission is quite important: people will be annoyed if there is too much of a delay in a voice connection. For data traffic, delay is less important, but the utilization of the available network bandwidth is more difficult to attain, because relative to voice, data traffic tends more “bursty” – the rate of information being sent varies very much over time.

Voice traffic has been traditionally sent on circuit switched networks, whereas data traffic has been traditionally sent on packet switched networks. The difference between circuit switched and packet switched networks is in their allocation policy of the bandwidth of the connection in the network. In circuit switched networks, a fixed fraction of the network is allocated to each connection, and that fraction belongs to the users throughout the duration of the connection. In packet switched networks, the amount of the bandwidth allocated to the connection varies depending upon the traffic pattern of the connection: the connection receives more bandwidth when it is actually sending data and less when it is not. In general, packet switched networks are better for traffic which is more bursty because they utilize the bandwidth in the network more efficiently, and circuit switched networks are better for traffic which is more time-critical (such as real-time voice or video) because circuit switched networks guarantee the full bandwidth required for the entire connection.

Because of the difference in the traffic patterns of voice and data, ISDN was originally proposed as two networks, one of which was circuit switched for voice, and one of which was packet switched for data. In [17], Turner describes a network architecture called *Fast Packet Switching* (FPS) which is completely packet switched, but which can effectively carry voice traffic as well as data traffic. There are several reasons why voice and data should be carried on a single FPS network:

- Shared resources allows systems to amortize costs over a wider customer base, and smooth demand.
- Voice bandwidth requirements can be reduced by a combination of compression and multiplexing of voice streams – i.e. voice is also somewhat bursty, and can take advantage of the efficiencies of packet switching.

- FPS allows the easy incorporation of new applications which can use voice, video and data connectivity at the same time.
- FPS is more flexible in bridging networks of different speeds, and consequently, allows upgrading of parts of the network while maintaining availability.

In the next section, we describe briefly the design principles behind Fast Packet Switching.

1.3 Fast Packet Switching

In order for packet switching to meet the performance requirements required to replace and supersede the circuit switched networks for voice, Turner proposed several design principles for FPS:

1. Move most of the responsibility for error detection, error correction and flow control from the link level up to higher level protocols which operate on an end-to-end basis.
2. To meet latency requirements, drop packets when queues in switches overflow – however, engineer the switches so that such packet losses occur extremely rarely. This technique is called *statistical multiplexing*.
3. To support statistical multiplexing, have most communication be connection-based, so that the network only allows connections when it believes it has the resources to handle the connections without dropping too many packets because of queue overflow.
4. Do not tie the standard for FPS to a particular performance level.
5. Expend equivalent engineering effort and expense – although it is clear today that packet switched networks are capable of high-performance, at the time of the formulation of the principles of FPS, packet switched networks (generally used to network computers) had much lower performance than circuit switched networks (used for the phone system).

1.4 Asynchronous Transfer Mode vs. Packet Transfer Mode

Turner's proposal for an Integrated Services Packet Network based upon FPS has evolved into several competing proposals for Integrated Networks. In this paper, we divide these proposals into two groups, which we loosely call Asynchronous Transfer Mode (ATM) and Packet Transfer Mode (PTM). ATM is a standard which has been developed by CCITT, and which has become the dominant manifestation of FPS [9] [16]. "PTM" can be used to describe a few network designs, such as IBM's plaNET, as well as the Frame Relay standard for data networks, which has also been standardized by CCITT [6].

The primary difference between ATM and PTM is that ATM packets are fixed-sized and small, and PTM packets are variable-sized and may be fairly large. There are other differences between the two, such as the routing methods, which we will not explore in this paper. *In this paper, we will compare the effectiveness of fixed-sized packets versus variable-sized packets in FPS networks.*

2 Comparison of ATM and PTM

ATM packets are fixed-sized packets with data fields of 48 bytes. Five bytes are used for header, so the length of the whole packet is 53 bytes. For the remainder of this paper, we will call ATM "packets" *cells*.

PTM packets are variable-sized packets, and in this paper, we consider packets with maximum data field of 4096 bytes – most designs for PTM networks have maximum data fields of 2048 to 8192 bytes, and we chose 4096 as a compromise. As in [6] and [12], we assume PTM packets with header lengths of 12 bytes.

Neither ATM nor PTM packets directly implement user-level data packets (such as IP packets) which we will call *frames*. In general, frames must be segmented at the source into either cells or packets, and then reassembled at the destination, and this work is performed in a higher level protocol. Since ATM cells are much smaller than the maximum PTM packets, in general, frames must be segmented into many more ATM cells than PTM packets.

2.1 Format Efficiency

The different formats of ATM and PTM have an effect on the amount of overhead which is incurred. There are five principal sources of overhead:

1. User data must be split into chunks in both the ATM and PTM formats, but since PTM packets can be 4096 bytes long, there is less overhead incurred by headers for the PTM format, since there will be less headers for PTM than for ATM for the same user frame.
2. Because ATM is a fixed-length format, user level frames may not be an integral multiple of the length of the data field for ATM. In that case, the rest of the data in the last ATM cell does not contain any useful information, but must be transmitted anyway.
3. Since PTM is a variable-length packet format, extra information must be inserted into the packet to distinguish the data within the packet from the marker for the end of the packet. This overhead is called “bit-stuffing”, and according to [12], this is a 3.2% overhead for the data portion of the PTM packet.
4. There are standard protocols (ATM Adaptation Layers) for using ATM cells to transmit variable length user frames, and these protocols consumes four bytes of the data field for additional protocol information. This effectively increases the header length to 9 bytes, and decreases the data length to 44 bytes. Note that this overhead is not present for applications which do not require variable length user frames, such as voice.

Figure 1 shows a graph comparing the efficiency of ATM cells versus PTM packets as a function of the length of the frame being sent. The efficiency of ATM cells (E_{ATM}) as a function of the length X of the user frame can be represented by the formula:

$$E_{ATM}(X) = \frac{X}{\lceil \frac{X}{44} \rceil (44 + 9)}$$

The efficiency of PTM packets (E_{PTM}) as a function of the length X of the user frame the can be represented by the formula:

$$E_{PTM}(X) = \frac{X}{12 \cdot \lceil \frac{X}{4096} \rceil + 1.032 \cdot X}$$

The “jagged” shape of the ATM curve is a result of the wasted information in the last ATM packet representing a frame – if the frame length is an integral multiple of the length of the ATM data field,




Figure 1: Format efficiency of variable-sized packets versus fixed-sized packets as a function of frame length. This graph is a modification of one from *Asynchronous Transfer Mode, Solution for Broadband ISDN* 1993, De Prycker

then the efficiency is at a maximum for ATM, but if the frame is just one byte more, then another ATM cell must be sent, which causes a large drop in the efficiency.

The maximum achievable format efficiency for ATM (including the AAL overhead) is about 83%, and the maximum achievable format efficiency for PTM is about 96%.

If the probability density function for the length of a user frame is represented by the function $P_{FL}(X)$, then the overall format efficiency is the following:

$$\text{format efficiency} = \sum_{l=0}^{\infty} P_{FL}(l)E(l)$$

This equation shows that the overall format efficiency does depend upon $P(X)$, the distribution of the frame lengths. *However, in general, PTM has a better format efficiency than ATM.*

2.2 Maximum Packet/Cell Data Size

It is often stated that the short cell length of ATM is inefficient for data applications, with the implication that larger cell lengths would have a better format efficiency for those applications. This is

somewhat misleading – because ATM has a fixed-length cell format, there are two main contributions to the overhead: the header to data ratio, and the data to frame length. The header to data ratio is represented in figure 1 by the line showing the maximum efficiency for ATM, and the data to frame length ratio is represented by the difference between the maximum ATM efficiency and the actual ATM efficiency.

As frame lengths become large, the second ratio becomes less important, but many data applications send a significant number of small frames – this is discussed later in this paper. Larger cell sizes mean that the ATM efficiency curve becomes even more jagged, and especially in the case of small frames (less than 64 bytes), this can have a severe impact on efficiency.

2.2.1 Voice Packetization and ATM cell lengths

ATM cells are short because of the desire to efficiently carry voice traffic. Voice packets are generally of uniform length and are short – the reason for this is because of the delay for collecting the voice samples. Voice is usually sampled with 8 bit samples at 8 KHz. 48 voice samples can fit into one ATM cell, and that takes 6 milliseconds to create. If the ATM cell were longer, then filling up the entire cell would require more time. If there is a long delay between the utterance and the reception of the utterance, the quality of the conversation is degraded considerably, and often, echo suppression hardware must be deployed.

When using ATM, *every* voice frame is 48 bytes long, and since voice does not require the additional 4 bytes of higher protocol information, a voice packet will achieve the highest possible data format utilization for ATM. Currently, the telephone companies carry more traffic and make more money than all of the other potential users put together – although it is possible that larger frames might only be partially filled in for voice, they would probably not be happy with this compromise to their format efficiency.

2.3 Packet/Cell Loss

In FPS networks, these are three sources of packet or cell loss [18]:

- *losses due to the transmission errors.* Bertsekas notes in [2] that losses due to transmission errors are not a first order effect for most networks implemented with modern technology, especially optic fiber networks.¹
- *losses due to queue overflow.* These losses can be controlled by selecting an appropriate length queue – longer queues will reduce losses.
- *losses due to excessive delay.* These losses can be controlled by controlling the queue length, and only are relevant in the case of real-time applications such as voice or video. Shorter queue lengths will reduce these losses. We do not examine this source of losses in this paper.

Packet/cell loss rate should be viewed from the viewpoint of the application using the network. Certain applications have a high tolerance for packet loss, such as real-time voice or video. If a packet or cell is lost, that will cause a momentary drop in the quality of the service, but there are simple mechanisms to handle such loss gracefully.

¹“For most links in use (with the notable exception of radio links), the probability of [transmission] error on reasonable-sized frames is on the order of 10^{-4} or less, so that this effect is typically less important . . . unfortunately, there are many analyses of optimal maximum frame length in the literature that focus only on this effect. Thus, these analyses are relevant only in those special cases where error probabilities are very high.” *Data Networks*, p97




Figure 2: Contributions to Frame Latency

For data applications, where the data must be received without any errors, the relevant measure of interest is the user-level *frame* loss rate. This *frame* loss rate has different characteristics depending upon whether the network is ATM or PTM. The loss of a single packet/cell of a frame in transit usually means the loss of the entire frame [3].

In [6], Cidon *et al* describe what they call an “avalanche” effect, which refers to the loss of an entire user frame by the loss of a single ATM cell. Since ATM cells in a queue are less correlated to the same user-level frame in comparison to PTM packets, the loss of several consecutive ATM cells will more likely mean the loss of several different frames in comparison to an equivalent data loss rate of PTM packets.

2.4 Utilization of Available Bandwidth

The utilization of the output bandwidth of the switches is one of the primary performance measures which we consider in comparing PTM with ATM. ATM has a big disadvantage compared to PTM because ATM must transmit almost 20% more bits to send the same amount of user data as PTM. However, ATM has an advantage over PTM because the buffer utilization is smoothed out due to the segmentation and multiplexing of ATM cells. ATM’s smaller cells require less queue memory, and therefore reduce the cell loss rate. As is argued in [12], in some cases ATM can overcome its format inefficiency by using less queue memory, and therefore perform at a higher effective bandwidth for an equal frame loss rate.

2.5 Latency

Figure 2 shows the contributions to the total latency of communications as seen by the user. The primary components of latency are the packetization and depacketization, the transmission delay, the switching delay, and the queueing delay. In general, the packetization and depacketization delay are dependent upon the application – in the case of voice, this delay can be a significant contribution to the total delay. In the case of data applications, this delay is not very significant because there are no real-time constraints which slow down the packetization.

The transmission delay is dependent upon bandwidth, geography and the speed of light, and the last two are largely beyond the scope of the engineer designing the system. The actual switching delay within the switches is not very high, since these switches can run at high speeds, but the queueing delays seen by packets can add a significant amount to the total latency of the packet.

The only leverage point in reducing the latency experienced by the user is in trying to reduce the time spent in the queue. Packetization/depacketization and transmission are largely fixed, and the switching delays are minimal – latency is controlled by controlling the number and bandwidth of connections and by adjusting the queue length.

2.5.1 Multiplexing and Pipeline Effects

Given the same bandwidth utilization, smaller packets will incur less latency than larger packets and identically sized packets will incur less latency than varying sized packets [18] [2].

There are two factors which argue for smaller packets. First, small packets allow the pipelining of frames through switches, which reduces the queue memory requirements and latency. PTM cells need to be fully received before being forwarded across the switch (store and forward).

The second factor is that ATM cells can be more easily multiplexed – however, since ATM cells are generated from user frames, the distribution of ATM cell generation will show some correlation, or “batchiness. The degree of batchiness is dependent upon how much the inputs are multiplexed, which depends upon what the ratio of the bandwidth of the switch inputs and the aggregate bandwidth of the switch outputs. If the inputs are much slower than the total output, then the traffic arrival will look more like a Poisson process, and therefore have lower latency. If the inputs are of similar speed as the output, then packets will arrive at the switch in about the same way that PTM packets will.

For PTM, multiplexing effects are not important if the average frame size is less than the maximum size of the PTM network. In general, PTM switches cannot (and do not) pipeline the variable sized packets (such as wormhole routing in some networks for parallel computers), because of the need for error-checking and the difficulty of design. If the average frame size is less than the maximum data payload of the PTM network, then the traffic pattern does not change because of multiplexing. In the cases we consider, multiplexing does not effect PTM.

2.5.2 Measures of Latency

Queue latency is often expressed in terms of average memory requirements, and can be translated to time units by dividing by the output bandwidth of the switch. Latency is related to packet loss probability because longer latencies imply more queue memory requirements, which implies higher probability of queue overflow and packet loss.

2.6 Traffic Patterns

It is difficult to make engineering decisions for the design of a network when it is uncertain what kind of traffic it is being built for. Indeed, FPS networks must be built to handle any kind of traffic because the applications which are dominant today may not be the applications which will be popular once the infrastructure is in place.

We have already discussed the influence of the telephone companies – since the vast majority of their current traffic is voice, they envision some sort of similar traffic patterns in the future, and they have used their influence to choose the current cell size for ATM.

Some work has examined the characteristics of data traffic for computer networks, both Wide Area and Local Area. Cáceres [4] [5] has found that TCP frame lengths are bimodal, with one peak at under 10 bytes, and one peak at around 512 and 536 bytes. Amer, *et al* [1] find similar a similar bimodal distribution, with one peak for frames less than 64 bytes and another peak for frames between 800 and 1200 bytes. Schmidt and Campbell [15] find one peak at under 64 bytes and another peak at around 500 bytes. Gusella [11] finds the same peak at under 64 bytes, and another peak at 1072 and 1500 bytes.

The empirical computer network studies for the most part show that most of the frames are either used for control, which are very small, or else very large data frames corresponding to the largest frame size in the protocol being considered.

Of course, those studies can only describe the type of traffic which exists today – future applications may have very different characteristics. Some work has postulated that most of the bandwidth will be taken up by video applications [8] and other “multi-media” applications which will mostly consist of very large frames. In reality, no one knows what the traffic patterns will look like.

2.7 Interaction with Higher-Level Protocols

FPS networks will be used to transport data using existing protocols and interfacing to existing, non-FPS networks. [10] argues that PTM more closely maps to the current protocols (TCP/IP) than ATM, because of its those protocols also have variable-sized frames. Cáceres [5] and Schmidt [15] both examine the interaction of the ATM standard with TCP/IP, and to an extent, both conclude that ATM does not map efficiently onto the higher level protocol, because of cell size decisions.

Other work [7] argues that new network architectures such as ATM require new network protocols to use them efficiently. TCP/IP was designed at a time when network capabilities and architectures were quite different than today, and such an argument may have merit. However, because of inertia, and because of the large installed user-base, such a conversion to newer protocols would be unlikely in the near term.

2.7.1 Segmentation and Reassembly Costs

Because of ATM's shorter maximum packet length, it will be more expensive to segment the protocol data unit level frames before injecting them into the ATM network, and then re-assembling them on the destination side. Cidon, Gopal, *et al* [10] argue that such costs would require extra computer power to handle the segmentation and reassembly costs. In our opinion, this work can be and will be handled by special-purpose hardware at the interface-level – note that not all applications will require such hardware, and therefore, such hardware cost should be shouldered by those applications which require them. For instance, telephone service does not require segmentation and reassembly, so telephones should not require special hardware, whereas computers will, and their interfaces to the network will provide such hardware.

2.8 Simplicity of Implementation

Perhaps one of the most important arguments for ATM is that its fixed-size cell format allows for a much simpler implementation. Variable-sized packets may require some sort of memory management for queues, and the variable-sized packets do not lend themselves to rapid dispatching. Fixed-sized cells make pipelining much simpler; however, the small size of the ATM cells requires rapid processing of headers for each cell.

Like the argument for RISC architectures, it is difficult to quantify the engineering effort differences. By considering the RISC experience however, it is clear that simpler designs allow higher clock speeds, require less logic, and take less design time, which accelerates the rate of innovation. We have already noted that switching time is not a significant contributor to total latency, but the flexibility afforded by ATM will allow larger and fancier switches to be built.

We will not attempt to quantify the benefits in regards to simplicity of implementation in this paper, but we believe it is an important factor.

2.9 Too Many Issues!!

There are many issues which must we must consider. In the next section, we give a system-level overview of some of the issues we discussed in this section, and describe a simple way to think about the relationship between these issues and characterize the arguments (pro and con) for ATM and PTM.

3 System Level Overview

3.1 Dependency diagram

Figure 3 shows a dependency diagram illustrating the dependencies between the variables and constants used in evaluating the performance of ATM networks. The arcs show a dependency, and the “+” and “-” signs on the arcs describe whether there is a positive or negative relation between the two variables. A “+” indicates that if the source variable increases, the destination variable increases, and if the source variable decreases, the destination variable decreases. A “-” indicates that if the source variable increases, the destination variable decreases, and if the source variable decreases, the destination variable increases.

The meaning of “decrease” and “increase” depends upon the variable – for instance, increasing *average cell latency* is bad, but increasing *user effective bandwidth* is good. All of the relationships between variables connected by an arc are monotonic – either monotonically increasing or monotonically decreasing. However, this diagram does not describe the specific relationship between the variables, and in almost all cases, these relationships are highly non-linear.

Some simplifications are made to make it easier to think about the relationships. The following are the definitions of the variables.

cell header The length of the cell header. This is 9 bytes for ATM, and is a constant.

cell payload The length of the cell data. This is 44 bytes for ATM, and is a constant.

cell length The length of the cell, which is 53 bytes for ATM, and this is a constant.

mean frame length The mean frame length. We ignore the shape of the frame length distribution to make the model simpler to think about. We do not know what the mean frame length of representative traffic of the future is, so this is a variable.

mean frame interarrival time The mean time between arrivals of user-level frames – again, we ignore the shape of the distribution to make the model simpler. This is a variable.

format efficiency The percentage of the total information transmitted which is what the user actually cares about. This was discussed in the previous section, and ATM has a lower format efficiency than PTM.




Figure 3: System level overview of dependencies between engineering decisions and performance

input multiplexing ratio This is the ratio of the input bandwidths to the aggregate output bandwidth. For example, a switch may have an aggregate output bandwidth of 100 Mbits/s, but be fed by 100 inputs which are each 1 Mbit/s – in this case, the multiplexing ratio is 100.

switch queue length The length of the queue in the switch – this is an engineering decision, so this is a variable, and this will effect the switch bandwidth utilization, the cell loss ratio and the average cell latency, as is shown in the diagram.

switch bandwidth utilization This is the utilization rate of the switch output, which is usually called ρ in queueing theory. this is dependent upon the mean frame length and the mean frame interarrival time. However, an inefficient format will, in effect, increase the the mean frame length or decrease the frame interarrival time. Longer queue lengths will also mean less cell loss, which maintains switch bandwidth utilization, so the switch bandwidth utilization is also a function of these two factors.

average cell latency The average latency seen by a cell – according to queueing theory, this increases as the bandwidth utilization increases, and as the cell length increases. Also, shorter queue length will decrease the latency by dropping more packets, and a higher input multiplexing ratio will decrease the average latency by smoothing out traffic demands.

cell loss ratio This is the percentage of cells which are lost, and is a variable of the average latency

and average queue length. Longer queues mean less cells are lost, and longer latencies mean more cells are lost.

frame loss ratio This is the percentage of *user* frames which are lost – this is the variable which the user is really concerned about, and it is a function of the input multiplexing ratio and average cell latency. A higher input multiplexing ratio will increase the frame loss ratio because successive cells in the queue will become less correlated with regard to their original user-level frames, and more frames will be lost to the “avalanche” effect described in [10]. A higher cell loss ratio will also increase the frame loss ratio.

switch total bandwidth This is a constant which describes the total output bandwidth.

user effective bandwidth This is the bandwidth which is available to the users of the system. This is dependent upon the actual switch bandwidth utilization and the efficiency of the cell format, as well as the frame loss ratio. A higher frame loss ratio will decrease the effective useable bandwidth, and a higher format efficiency will increase the user effective bandwidth.

3.2 How to think about ATM versus PTM arguments

Using this diagram, we can evaluate some of the arguments which are given in the debate between ATM and PTM. For example, Cidon *et al* argue that the ATM cell latencies are actually *not* lower than PTM cell latencies, and they give two reasons. One of which is that the *input multiplexing ratio* is actually rather low, and successive cells will likely be correlated and arrive in bursts, leading to traffic patterns similar to PTM. i.e. ATM will have the same arrival characteristics as PTM, except for the added overhead of the cell format. The second reason they give is that *format efficiency* for ATM is low (which we saw in the last section) and that *switch bandwidth utilization* is therefore high, therefore increasing average cell latency. We will see some evidence of the second reason in our simulation study in the next section.

Cidon *et al* then argue that *frame loss ratio* is high in ATM because the *input multiplexing ratio* is high, which mixes up cells representing different user frames, and therefore leads to higher frame loss due to what they call the “avalanche effect” – one lost cell can lose the entire frame, and cells are not highly correlated. This argument is represented by the one arc between *input multiplexing ratio* and *frame loss ratio*, and it somewhat contradicts their assertion that successive cells will be correlated in ATM and therefore lead to high latencies. Furthermore, they do not note the path between *input multiplexing ratio*, *average cell latency*, *cell loss ratio* and *frame loss ratio*. This is described by Le Boudec in [12], where he notes that the increase in frame loss due to the “avalanche effect” is sometimes made up by a decrease in the *average cell latency*, causing a decrease in the *cell loss ratio*, causing a decrease in the *frame loss ratio*. Again, we will see some evidence of the effect described by Le Boudec in the simulation studies at very high input multiplexing ratios.

We do not draw a similar diagram to represent PTM, because most of the diagram will be the same, there will be no *cell length* variable – there will simply be an arc between the *mean frame length* and the *average cell latency*. There will also be no *input multiplexing ratio* variable. The *format efficiency* will be higher than for ATM.

3.3 Performance variables are function of engineering decisions

The three variables which the user is concerned about are at the bottom of the system-level overview diagram (Figure 3), and they are the *user effective bandwidth*, the *frame loss ratio* and the *average cell latency*. The second and the third performance measure have upper bounds, but performance




Figure 4: Overview of Network Simulator

which exceeds the upper bounds does not necessarily significantly affect the end user, for most current applications. The first performance figure is important to the operators of the network, because it refers to the number and bandwidth of the connections which can be made in the network. The higher the total bandwidth, the more money they can make, and in an indirect way, this is a figure of interest to the end user because it may affect his ability to make a connection.

Some of the variables are actually constants, such as the cell length and switch total bandwidth, and we do not discuss the possibility of changing these constants – that is not the purpose of this study. Some of the variables are beyond the scope of the engineer, but uncertain, such as the mean frame length and mean frame interarrival time. Some variables are the ones the engineer must adjust to maximize performance (by meeting a latency requirement and a frame loss ratio requirement).

In the next section, we describe a simulator we use to explore the space of the variables which we can alter, and their effect on the variables which are the performance measures we will be concerned with.

4 Simulation Results

To evaluate the effect of the variable factors described in the previous section on the performance measures of interest, we have built a simple FPS network switch simulator. Analytical models are interesting and useful, but are often quite cumbersome to develop and limited in their applicability. For instance, Le Boudec [12] uses two different analytical models to describe ATM behavior, depending upon the length of the queue buffer, and both models were inadequate to describe even longer queues. Cidon *et al* [10] use an ATM latency model described by Parekh and Sohraby [14] to argue for PTM, but Parekh and Sohraby use simulations to present their results. Naghshineh and Guerin [13] use simulation to analyze queue buffer usage and error rates as a function of utilization and multiplexing, which is similar to our goals.

The results from our simulations are not surprising, and in fact agree qualitatively with the analytical and simulation results of the previous authors described. Some of the actual graphs may appear slightly different because of different assumptions about the format overheads and multiplexing ratios. It may appear odd that two authors who come to opposite conclusions (Cidon and Le Boudec) as well as two authors who conclude that performance cannot be a motivating factor to decide on PTM versus ATM (Parekh and Naghshineh) all agree. In fact, they do, and some of the differences are the result of different initial assumptions, and some of the differences are the result of seeing the glass half empty or half full.

4.1 Simulator Overview

Figure 4 shows a high-level overview of the components of the switch simulator. There are three basic stages in the simulator: the raw trace generator, the PTM or ATM trace generator, and the queue simulation. Each of these parts can be easily modified to change the model if necessary, and the output of the raw trace generator can also be fed directly into the queue simulation.

4.1.1 Raw Trace Generator

The raw trace generator creates a stream of user-level frames, each of which is represented by a time-stamp and a frame-length. The generator itself can create the trace in a number of ways, but for this study we generated traces with Poisson process arrival times and frame lengths which were exponentially distributed.

The Poisson arrival time assumption allows us to consider the one trace an interleaving of several slower sources with the same frame size distribution. In that way, we do not have to deal with separate traces for separate sources.

The aggregate user utilization of a trace is simply the quotient of the mean frame length by the mean interarrival time divided by the aggregate output bandwidth.

4.1.2 ATM and PTM Trace Generator

ATM and PTM traces can be generated from raw traces, and they have the same data representation as raw traces, except that each packet or cell in the trace will be tagged with an identifier describing the original frame which it represents. The segmentation of the original frame into ATM and PTM packets is performed with the format overhead assumptions described earlier in the paper. Each ATM cell will require 9 bytes of header and can only carry 44 bytes of data payload. Each PTM packet has a 12 byte header, and the length of the user frame is expanded by 2.3% to account for stuffing – we assume a maximum PTM packet payload of 4096 bytes.

The interarrival times of the cells or packets corresponding to the same user-level frame are determined by a “multiplexing” factor. The higher the multiplexing factor, the longer the interarrival time of the cells corresponding to the same frame. The multiplexing factor can be considered the ratio between the aggregate output bandwidth of the switch and the bandwidth of a single input.

The output utilization characteristics of ATM and PTM traces generated from the same raw trace will be different because ATM has a higher format overhead than PTM. The ATM trace will have a utilization rate which is about 20% higher than the corresponding PTM trace. In all of the simulations we run, we will always describe the ATM and PTM traces according to the utilization of the original user trace.

4.1.3 Single Server Queue Simulator

The queue simulator itself is a single server model. The simulation itself is performed using time-warping on the trace packets and each packet is assumed to take up as much queue space as the length of the packet for a duration which is proportional to the length of the packet plus the time it takes for any packets already in the queue to leave. If this is combined with a raw trace which is Poisson-distributed along the arrival times, and exponentially distributed for the frame lengths, then this will simulate an M/M/1 queue.

The simulator does not simulate output contention in detail, because all of the packets passing through the switch are aggregated into the same queue. Although this is a very crude model, it is in essence identical to the assumptions made in the studies comparing fixed-length and variable-length packet formats [10] [12] [14] [13].




Figure 5: Comparison of simulator and queueing theory results for an M/M/1 queueing system and an M/D/1 queueing system

4.1.4 Simulator Test

To test out the accuracy of the simulator, we have simulated two well-known queueing systems and compared them against the theoretical behavior. The M/M/1 queueing system models Poisson arrivals and exponentially distributed frame lengths. The M/D/1 queueing system models Poisson arrivals and fixed frame lengths.

We set a mean frame length of 700 bytes for both the M/M/1 and M/D/1 systems, and the results are shown in Figure 5. The simulated results conform quite well to the expected theoretical results, and furthermore, this graph demonstrates the advantage of fixed-sized packets over variable-sized packets, even when the mean packet lengths are the same. In the following sections, we describe some of the experiments we ran on the simulator and the implications of the results.

4.2 Average Latency versus User Utilization

Figure 6 shows the queue latency as a function of the user utilization for both ATM and PTM. This is a similar graph to the previous graph. The mean user frame length is 700 bytes, and the lengths are exponentially distributed.

The PTM curve looks similar to the latency versus user utilization curve shown in the test case for the M/M/1 case. That is because PTM is almost the same as a simulation of the raw M/M/1 queue – there is some small additional overhead for the bit-stuffing and the header, but it looks very similar to the M/M/1 case.

The ATM curves all seem to have an asymptote at around 80% user utilization. The reason for this is very easy to explain – a 80% user utilization is almost a 100% ATM utilization because of the additional overheads for the ATM headers, and as can be seen in the M/D/1 graph, the latency curve




Figure 6: Latency as a function of the User Utilization. Note that the ATM curves all have asymptotic behavior when the User Utilization Ratio reaches around 83% – this is because the actual utilization ratio being experienced is near 100% because of the additional overhead in the ATM format

increases exponentially as the utilization goes beyond 80% utilization. 80% utilization for ATM is approximately 66% user utilization – this is clearly seen in figure 6.

Note that as the multiplexing rate increases, the ATM latency curve decreases – this is the smoothing effect described earlier. When the multiplexing rate is not high, the ATM cell arrive in bunches, which makes the behavior more like M/M/1 with mean packet size equal to the user’s mean frame size plus the header overheads; when the multiplexing rate is high, the ATM cell appear less correlated, and the behavior is more like M/D/1 with mean packet size being 53 bytes, the size of the ATM cell.

Figure 6 indicates that ATM is better than PTM at low utilization ratios and high multiplexing ratios, but at higher utilization and lower multiplexing ratios, PTM is better.

4.3 Probability of Overflow versus Queue Length

Figure 7 shows the effect of user utilization on the queue overflow probability as a function of the queue length. Again, the mean frame size is 700 bytes, and for these two graphs, the multiplexing rate is 4. Most of the overflow curves eventually level out at a certain queue length, which indicates it is unlikely at that length that a queue will overflow, given a particular load.

At low user loads, ATM has an advantage over PTM, even though it has considerably more overhead – the multiplexing smooths out the demand on the buffer. At a user utilization level of about 0.7, the ATM and PTM curves appear to level out at about the same place. This is the place where the ATM curve begins to increase exponentially, and for the highest user utilization level, ATM has a much higher queue overflow probability than PTM.

4.4 Probability of Overflow versus Multiplexing




Figure 7: Probability of overflow as a function of queue size. Note that ATM uses much less buffer when the load is low, but uses much more buffer when the load becomes high.

Figure 8 shows the effect of multiplexing on the overflow probability for ATM networks only; multiplexing does not have an effect on PTM networks for the frame length distribution we are considering. The user utilization rate is .5, which is not in the exponential region for ATM, and the dashed line indicates the overflow rate for an idealized user load, without format overhead or segmentation.

Note that with low multiplexing, ATM actually has a higher overflow rate than the idealized user load – this is because only the bad effects of the header overhead are felt, whereas the beneficial effects of the segmentation to small cells are not seen until the multiplexing rate is reasonably high. As the multiplexing rate increases, the overflow rate decreases. For moderate rates of multiplexing (5X), we see that a buffer size of between 10,000 and 15,000 bytes are necessary to maintain cell losses of less than 10^{-5} . For higher rates of multiplexing, we can use much smaller queues to maintain the same loss ratio.

4.5 Frame Loss versus Multiplexing and User Utilization

Because the loss of a single ATM cell will usually cause the loss of an entire user frame, Cidon *et al* [10] contend that ATM will cause catastrophic losses because cells are not correlated, and thus a queue overflow will tend to cause the loss of many different frames. This implies that they believe higher multiplexing rates will cause higher frame loss. Le Boudec [12] claims that the increased multiplexing rates will decrease queue usage, and thus will decrease cell loss faster than the increase in the avalanche effect.

This graph shows that Le Boudec is correct: higher multiplexing rates improve frame loss rates because the avalanche effect is not as strong as the the improvement in the buffer usage caused by multiplexing.




Figure 8: Probability of overflow as a function of queue length, varying *multiplexing* rate. If an ATM switch is receiving its traffic from highly multiplexed inputs, the multiplexing will have effect of smoothing out the load, and reducing the spikes in buffer size caused by the entrance of large frames. If there is little multiplexing, then the ATM behavior is similar to PTM behavior, since cells will be closely spaced together and look similar to PTM packet arrivals.

4.6 Caveats and Future Work

Our simulation work makes several simplifying assumptions in order for the study to be tractable. In comparison to analytical work, it is simple to try different distributions for frame interarrival time and frame size as well as architectures. In addition, we have not considered the problem of congestion in routing which would require postulating about specific switch architectures – such work would not be as generally applicable as the simple model we have implemented.

The results are based on the assumptions we have made about header and data field sizes and other overheads. The qualitative nature of the results will still be relevant for small variations in these parameters, but some of the curves may be shifted over.

5 Proposed Experiment

Unfortunately, simulations and analysis are not enough to make conclusive statements – both simulation and analysis must make simplifying assumptions about network traffic, architecture, topology, queueing policies, and technology in order to be merely feasible. A real study would compare the performance of ATM and PTM on real traffic in a real environment on real hardware.

The plaNET project at IBM is perhaps the most serious implementation of a PTM network whose intended traffic is the same as for ATM networks. In [10], some of the designers of plaNET describe some of the recent modifications to plaNET which allow it to support ATM style traffic without overhead in terms of the packet format. The internal format of ATM-style packets in plaNET is not identical to ATM, but it is merely a rearrangements of the various fields of ATM. Routing for these packets is




Figure 9: Probability of User Frame loss as a function of the User Utilization. The ATM curves show the beneficial effects of multiplexing in lowering the queue usage outweigh the bad effects of increased “avalanche” effect.

performed in label swapping manner, just as in ATM, and plaNET clusters can act as bridges between ATM networks.

The plaNET network is the ideal environment to perform performance comparisons between ATM and PTM on real hardware with real traffic in a real environment. It supports both ATM and PTM, and it has the same parameters in terms of technology, queue length, topology, and architecture. Naturally, plaNET was designed primarily with support of PTM in mind, and therefore, may have some architectural decisions which may favor PTM – these factors should be taken into account in such a comparison.

6 Conclusion

In this paper, we have described the various arguments for ATM and PTM, and we have described a system-level framework to clarify the reasoning behind the individual arguments for each. To get a feel for the relationships between the possible variables which can affect the performance measurements for each, we built a simulator and performed experiments to show general characteristics of ATM and PTM switch behavior.

Our results qualitatively agree with the results of in [10] [12] [13] [14]. At low user utilizations, ATM is superior to PTM, but at higher utilizations, ATM reaches the saturation region (when the latency increases exponentially as a function of the bandwidth utilization) more quickly than PTM because of its higher format overheads. Such overheads will limit the bandwidth utilizations of ATM networks to be about 60-70% of user utilization – ATM itself reaches the saturation region when it is using 80% of the output bandwidth, but the extra overhead of ATM decreases the *user* utilization to about 65%.

The latencies and frame and cell loss rates in ATM networks are a function of the multiplexing ratio of the switch. Both the latencies and cell loss rates decrease significantly when multiplexing rates are

high, causing ATM to be superior to PTM when both are not in the saturation region. In addition, the “avalanche” effect caused by the loss of a user frame with the loss of a single ATM cell is overcome by the decrease in the cell loss rate when the multiplexing ratio increases.

Each network has its advantages, depending upon the demands of the applications. In the case of an applications such as voice which demand low latencies, and is characterized by a high multiplexing rate, ATM will be superior. In applications such as most computer data applications where latency is not as critical, but bandwidth utilization is very important, PTM will allow a higher utilization of the bandwidth of the channel, but at a cost of higher latency even when the network is relatively lightly loaded.

References

- [1] Paul D. Amer, Ram N. Kumar, Ruey-bin Kao, Jeffery T. Phillips, and Lillian N. Cassel. Local Area Broadcast Network Measurement: Traffic Characterization. In *Proceedings of Compcon '87*, 1987.
- [2] Dimitri Bertsekas and Robert Gallager. *Data Networks*. Prentice Hall, 2nd edition, 1992.
- [3] Andre B. Bondi and Wai-Sum Lai. The influence of cell loss patterns and overhead on retransmission choices in broadband ISDN. *Computer Networks and ISDN Systems*, 26(5):585–598, January 1994.
- [4] Ramón Cáceres. Measurements of Wide Area Internet Traffic. Technical Report TR-89-550, Berkeley, December 1989.
- [5] Ramón Cáceres. Efficiency of ATM Networks in Transporting Wide-Area Data Traffic. Technical Report TR-91-043, ICSI/Berkeley, July 1991. Obtained by anonymous FTP from datanet.tele.fi.
- [6] Israel Cidon, Jeff Derby, Inder Gopal, and Bharath Kadaba. A Critique of ATM from a Data Communications Perspective. *Journal of High Speed Networks*, 1(4):315–336, November 1992.
- [7] David D. Clark and David L. Tennenhouse. Architectural Considerations for a New Generation of Protocols. In *Sigcomm '90 Symposium on Communications Architectures and Protocols*, pages 200–208, September 1990.
- [8] Martin De Prycker. Definition of Network Options for the Belgian ATM Broadband Experiment. *IEEE Journal on Selected Areas in Communications*, 6(9):1538–1544, December 1988.
- [9] De Prycker, Martin. *Asynchronous Transfer Mode, Solution for Broadband ISDN*. Ellis Horwood, 1993.
- [10] Inder Gopal, Roch Guérin, Jim Janniello, and Vasilios Theoharakis. ATM Support in a Transparent Network. *IEEE Network*, 6(6):62–68, November 1992.
- [11] Riccardo Gusella. The Analysis of Diskless Workstation Traffic on an Ethernet. Technical Report TR-87-379, Berkeley, November 1987.
- [12] Jean-Yves Le Boudec. About Maximum Transfer Rates for Fast Packet Switching Networks. *ACM SIGCOMM, Computer Communication Review*, 21(4):295–304, September 1991.
- [13] Mahmoud Naghshineh and Roch Guérin. Fixed Versus Variable Packet Sizes in Fast Packet-Switched Networks. In *Proceedings of IEEE Infocom '93, Volume 1*, pages 2c.2.1–2c.2.10, March 1993.

- [14] Shyam Parekh and Khosrow Sohraby. Some Performance Trade-offs Associated with ATM Fixed-Length vs. Variable-Length Cell Formats. In *Proceedings of IEEE Globecom '93, Volume 3*, pages 39.4.1–39.4.6, 1988.
- [15] Andrew Schmidt and Roy Campbell. Internet Protocol Traffic Analysis with Applications for ATM Switch Design. Technical Report UILU-ENG-92-1715, University of Illinois, May 1992.
- [16] Telco Systems. Asynchronous Transfer Mode: Bandwidth for the Future.
- [17] Jonathan S. Turner. Design of an Integrated Services *Packet* Network. *IEEE Journal on Selected Areas in Communications*, SAC-4(8):1373–1379, November 1986.
- [18] Jonathan S. Turner and Leonard F. Wyatt. A Packet Network Architecture for Integrated Services. In *Proceedings of IEEE Globecom '83*, pages 45–50, 1983.

Fixed-Length Packets versus Variable-Length Packetsin Fast Packet Switching

NetworksAndrew Shaw · Published 1994 · Computer Science, Engineering

TLDR This paper evaluates one design parameter of FPS networks: whether the packets should be xed-length or variable-length, and considers three measures of performance: user frame loss rates, average latency, and eective bandwidth. [Expand](#)

No Paper Link Available

Save

Alert



18 References

Related Papers

18 References

Search authors, publi

Q

Date Range

Citation Type

Has PDF

Author

More Filters

More Filters

Filters

Sort by Most Influe...

Fixed versus variable packet sizes in fast packet-switched networks

M. Naghshineh R. Guérin Computer Science, Engineering · IEEE INFOCOM '93 The Conference on Computer... · 1993

TLDR These performance measures include queue length distribution, packet loss probability, and user frame loss probability are investigated, when fast packet-switched networks that operate with either fixed or variable packet sizes are compared. [Expand](#)

26

3 Excerpts

Save

About maximum transfer rates for fast packet switching networks

J. Le Boudec Computer Science, Engineering · SIGCOMM '91 · 1991

TLDR An analysis framework for comparing the maximum achievable data rates that can be expected when ATM cells or variable length packets are used for L.AN interconnection shows that in several cases the ATM can achieve a higher useful rate in spite of the overhead handicap. [Expand](#)

4

Highly Influential

PDF

14 Excerpts

Save

Internet protocol traffic analysis with applications for ATM switch design

Andrew Schmidt R. Campbell Computer Science, Engineering · CCRV · 1993

TLDR The behavior of local and wide area Internet Protocol networks is examined, connection and connectionless protocols are studied, and the results are used to estimate the behavior of a virtual circuit network replacement for these networks. [Expand](#)

35

Highly Influential

5 Excerpts

Save

Design of an integrated services packet network

J. Turner Computer Science, Engineering · SIGCOMM '85 · 1985

TLDR It is argued that the evolutionary approach inherent in current ISDN proposals is unlikely to provide an effective long term solution and advocated a more revolutionary approach, based on the use of advanced packet switching technology. [Expand](#)

216

Highly Influential

PDF

3 Excerpts

Save

Definition of network options for the Belgian ATM broadband experiment

M. Prycker Computer Science, Engineering · IEEE J. Sel. Areas Commun. · 1988

TLDR The author describes the options and parameters which have been selected for a Belgian experiment planned in the early 1990s, including fixed or variable packet length, the optimal packetlength, transmission and switching speed, and the number of lower-layer capabilities. [Expand](#)

23

1 Excerpt

Save

ATM support in a transparent network

I. Gopal R. Guérin J. Janniello Vasilios Theoharakis Computer Science, Engineering · IEEE Network · 1992

TLDR The approach taken and components used in plaNET to provide the simple interfacing with ATM (asynchronous transfer mode) networks and links are described.[Expand](#)

 13  Save

Some performance trade-offs associated with ATM fixed-length vs. variable-length cell formats

S. Parekh K. Sohraby Computer Science, Engineering · IEEE Global Telecommunications Conference and... · 1988

TLDR It is concluded that the performance differences found here do not overwhelmingly favor either an FLC or a VLC format, and thus it is important to factor in other issues as well.[Expand](#)

 9  Highly Influential  4 Excerpts  Save

The Analysis of Diskless Workstation Traffic on an Ethernet

R. Gusella Computer Science, Engineering · 1987

TLDR The study suggests that protocol behavior is sup-optimal on a medium- size Ethernet local area network that connects file servers to diskless workstations, and details the behavior of each of the protocols responsible for a significant proportion of the total traffic.[Expand](#)

 36  1 Excerpt  Save

Asynchronous Transfer Mode, Solution for Broadband Isdn

M. Prycker Computer Science, Engineering · 1991

A resource for the design or installation of ISDN that covers transfer mode, ATM switching, broadband transmission, broadband subscriber premises networks—including MANs (metropolitan area... [Expand](#)

 785  Highly Influential  19 Excerpts  Save

Architectural considerations for a new generation of protocols


D. Clark D. Tennenhouse Computer Science, Engineering · SIGCOMM '90 · 1990

TLDR This paper identifies two new design principles, Application Level Framing and Integrated Layer Processing, and identifies the presentation layer as a key aspect of overall protocol performance.[Expand](#)

 1,417  PDF  1 Excerpt  Save

 1 2 

Related Papers



Stay Connected With Semantic Scholar

 [Sign Up](#)

What Is Semantic Scholar?

Semantic Scholar is a free, AI-powered research tool for scientific literature, based at the Allen Institute for AI.

[Learn More](#)

About

- [About Us](#)
- [Meet the Team](#)
- [Publishers](#)
- [Blog](#)
- [AI2 Careers](#)

Product

- [Product Overview](#)
- [Semantic Reader](#)
- [Scholar's Hub](#)
- [Beta Program](#)
- [Release Notes](#)

API

- [API Overview](#)
- [API Tutorials](#)
- [API Documentation](#)
- [API Gallery](#)

Research

- [Publications](#)
- [Researchers](#)
- [Research Careers](#)
- [Prototypes](#)
- [Resources](#)

Help

- [FAQ](#)
- [Librarians](#)
- [Tutorials](#)
- [Contact](#)

